Semantic Web – Interoperability, Usability, Applicability 0 (0) 1
IOS Press

A Reasonable Semantic Web

Editor(s): Krzysztof Janowicz, Pennsylvania State University, USA
Solicited review(s): Claudia d'Amato, Università degli Studi di Bari, Italy; Thomas Lukasiewicz, University of Oxford, UK
Open review(s): Aidan Hogan, DERI Galway, Ireland; Axel Polleres, DERI Galway, Ireland

Pascal Hitzler [a], Frank van Harmelen [b]
[a] *Kno.e.sis Center, Wright State University, Dayton, Ohio, U.S.A.*
[b] *Vrije Universiteit Amsterdam, The Netherlands*

Abstract. The realization of Semantic Web reasoning is central to substantiating the Semantic Web vision. However, current mainstream research on this topic faces serious challenges, which forces us to question established lines of research and to rethink the underlying approaches. We argue that reasoning for the Semantic Web should be understood as "shared inference," which is not necessarily based on deductive methods. Model-theoretic semantics (and sound and complete reasoning based on it) functions as a gold standard, but applications dealing with large-scale and noisy data usually cannot afford the required runtimes. Approximate methods, including deductive ones, but also approaches based on entirely different methods like machine learning or nature-inspired computing need to be investigated, while quality assurance needs to be done in terms of precision and recall values (as in information retrieval) and not necessarily in terms of soundness and completeness of the underlying algorithms.

Keywords: Semantic Web, Formal Semantics, Knowledge Representation, Automated Reasoning, Linked Open Data

1. The Linked Data Web needs semantics

The Semantic Web community, in the course of its existence, has gone through an interesting swing concerning the emphasis between "data" and "knowledge."[1] Indeed, much of the talk (and research, and writing, and programming) in the early days of the Semantic Web was about ontologies as objects of study in their own right: languages to represent them, logics for reasoning with them, methods and tools to construct them, etc. Many of the research papers in the first half decade of Semantic Web research (say, 1999-2005) seemed to forget that ontologies are not made for their own sake, but that the purpose of an ontology (at least on the Semantic Web), is to help foster semantic interoperability between parties that want to exchange data. In other words, the knowledge in the ontologies (the T-box) is supposed to help interoperability of the data (the A-box).

This insight was at the birth of the Linked Open Data project [2], which put a renewed emphasis on publishing sets of actual data according to web principles. However, as it is often the case with "counter-movements," it seems to us that (some of) the Linked Open Data work is erring on the other side, by only publishing just the data, and ignoring the value that can be had by annotating the data with shared ontologies.

Some of the problems that are plaguing the current Linked Open Data sets can be profitably solved by annotating data with ontologies. For example, knowing that some properties are inverse functional, knowing that certain classes are contained in each other, or that other classes are disjoint, all help to solve the instance unification problem.[2]

Similar arguments have been put forth regarding querying of Linked Open Data [19]: One of the main obstacles in querying over multiple Linked Open Data datasets is that severe information integration issues require solving. While having all data in RDF syntax (Resource Description Framework [23]) solves the information integration issue on a syntactic level, the current state of querying over the Linked Open Data cloud exposes the fact that semantic integration is hardly present. Indeed, RDF language primitives which are actually reflected by the RDF formal semantics (such as `rdfs:subClassOf` or

[1] or, in Description Logic speak: between "A-box" and "T-box"

[2] The instance unification problem refers to the problem of determining when two differently named instances are in fact identical.

```
bills/h3962       dc:title          "H.R. 3962: ..." ;
                  usbill:hasAction _:bnode0 .
_:bnode0          usbill:vote       votes/2009-887 .
votes/2009-887    vote:hasOption    votes/2009-887/+ .
                  dc:title      "On Passage: H.R. 3962 ..." ;
votes/2009-887/+ rdfs:label         "Aye" ;
                  vote:votedBy      people/P000197 .
people/P000197    usgovt:name       "Nancy Pelosi" .
```

Fig. 1. GovTrack triples encoding the knowledge that Nancy Pelosi voted in favor of the Health Care Bill. URIs have been abbreviated freely since the details do not matter for our discussion.

`rdfs:domain`) are relatively scarce in the cloud.[3] The only strong semantic language primitive used heavily is `owl:sameAs` from the Web Ontology Language OWL [15], and it has been observed frequently that its use is often rather abuse [6,13].

Another issue which points at a lack of semantics is the sometimes rather convoluted way of expressing knowledge in the Linked Open Data cloud. As just one example, let it be noted that the simple fact *Nancy Pelosi voted in favor of the Health Care Bill* is encoded in GovTrack[4] using eight RDF triples, two of which share a blank node (see Figure 1). From this and other examples, it seems apparent that triplification for the Linked Open Data cloud is often done without deep contemplation of semantic issues,[5] or of usefulness of the resulting data.[6]

2. Semantics as shared inference

Semantic interoperability is usually defined in terms of a formal semantics. But what does it mean for two agents to agree on the formal semantics of a message? Although the primary definition of the semantics of formal languages is most often in terms of a denotational semantics, e.g. [14] and [24] for RDF and OWL, respectively, perhaps a more productive definition on the Semantic Web is to describe semantic interoperability in terms of shared inferences.

When an agent (a web server, a web service, a database, a human in a dialogue) utters a message, the message will often contain more meaning than only the tokens that are explicitly present in the message itself. Instead, when uttering the message, the agent has in mind a number of "unspoken," implicit consequences of that message. When a web page contains the message "Amsterdam is the capital of The Netherlands," then some of the unspoken, implicit consequences of this are that Amsterdam is apparently a city (since capitals are cities), that The Hague is not the capital of the Netherlands (since every country only has precisely one capital), that The Netherlands is a country, or a province, but not another city, since countries and provinces have capitals, but cities do not; a spatial implied fact is that the location of the capital city is inside the area covered by the country, etc.

If agent A utters the statement about Amsterdam to agent B, they can only be said to be truly semantically interoperating if B not only knows the literal content of the phrase uttered by A, but also understands a multitude of implicit consequences of that statement which are then shared by A and B. It is exactly these shared, implicit consequences which are made explicit in the form of a shared ontology.

We could say that the amount of semantic interoperability between A and B is measured by the number of new facts that they both subscribe to after having exchanged a given sentence: the larger and richer their shared inferences, the more semantically interoperable they are.[7]

A language such as RDF Schema [23] which contains (almost) no negation, allows agent A to enforce beliefs on the receiving agent B, e.g. by specifying the domain and range of a property like "is capital of." This puts a *lower bound* on the inferences to be made by agent B, i.e., it "enforces" inferences to be made by B when it subscribes to the shared semantics. A richer language such as OWL [15] also allows agent A to "forbid" agent B to make certain inferences. Stating that Amsterdam is the capital of The Netherlands, that "is capital of" is an inverse functional property, and that Amsterdam is different from The Hague will disallow the inference that The Hague is the capital of The Netherlands. This puts an *upper bound* on the inferences to be made by agent B. By making an ever richer ontology, we can move the upper and lower bounds of the shared inferences ever closer, hence obtaining ever finer-grained semantic interoperability through an ever more precisely defined set of shared inferences.

Of course, this perspective of semantics as "shared inference" is entirely compatible with the classical view of semantics as model theory, in the sense of the formal semantics of, e.g., RDF and OWL: Valid inferences are inferences which hold in all models,

[3] "Scarcity," in this case, is a rather subjective matter. Let's just say that it currently seems too scarce to be really useful for reasoning.

[4] `http://www.govtrack.us/`

[5] See also [1,17,28] for further discussions.

[6] For an amusing critique on this practice, see [35].

[7] Ontology alignment issues obviously occur here, too.

and invalid inferences are inferences that hold in no model. However, semantics as "shared inference" does not presuppose the use of model theory,[8] although the latter currently seems to be the most advanced method for capturing this kind of semantics. Essential to the "shared inference" perspective is that it facilitates communication (and, thereby, interoperability), while model theory is often construed[9] as "the defining of meaning in a unique way."

3. Semantics as a gold standard

The usual role of semantics is to define precisely how the meaning of a set of sentences in a logic is defined. In Section 2, we have already seen that it is also possible to think of semantics in terms of an ever narrowing gap of multi-interpretability (with an ever increasing set of axioms closing the gap between what must be derived (inferential lower bound) and what may not be derived (inferential upper bound) from a set of sentences.

The classical view on semantics is then that any properly defined system must precisely obey this semantics: it must be sound and complete, i.e., any consequence prescribed by the semantics must also be derived by the system, and vice versa. Only recently the semantic web community has begun to appreciate the value of incomplete systems [11]. It is often useful to build systems that do not manage to derive all required consequences, as long as they derive a useful subset of these.

Rather than regarding this as an unfortunate but perhaps inevitable sloppiness of such implementations with respect to their semantic specification, we would advocate a different perspective, namely to view the formal semantics of a system (in whatever form it is specified) as a "gold standard," that need not necessarily be obtained in a system (or even be obtainable). What is required from systems is not a proof *that* they satisfy this gold standard, but rather a precise description of the *extent to which* they satisfy this gold standard [29].

Notice that in other, related, fields this is already commonplace: in Information Retrieval, the measures of precision and recall correspond exactly to soundness and completeness, but with the crucial difference that nobody only expects systems where both of these values are at 100%. Instead, systems are routinely measured on the extent to which they approximate full precision (soundness) and recall (completeness), and both researchers and application builders have learned to live with imperfect systems, and with laws that tell us that increasing one of the measures typically decreases the other. In short, the logical model has perhaps confused the ideal with the realistic, and the theory and practice of information retrieval may well be more appropriate for Semantic Web reasoners.[10]

A wide misconception is that, even when incompleteness may be a worthy strategy, surely unsoundness is bad in all cases. Again, the perspective from Information Retrieval shows that this is simply false: depending on the use-case, one may have a preference for erring either on the side of incompleteness (e.g. finding just a few but not all matching products is fine as long as all answers do match the stated requirements) or on the side of unsoundness (e.g. finding all potential terrorist suspects, even when this possibly includes a few innocent people). Just as in Information Retrieval, a use-case specific balance will have to be struck between the two ends of the spectrum, with neither being always better than the other.

From this perspective (semantics as a, possibly unobtainable, gold-standard) systems with anytime behaviour also become a very natural object of study: they just happen to be systems that succeed in increasingly better approximations of the gold standard as time progresses. It turns out that many algorithms for deduction, query answering, subsumption checking, etc., have a natural anytime behaviour that can be fruitfully exploited from the perspective of "semantics as a gold standard" that need not be perfectly achieved before a system is useful.

4. Semantics as possibly non-classical

If we take the viewpoints that "semantics is a (possibly unobtainable) gold standard for shared inference,"

[8] We do not want to propose any particular approach at this stage, but let it be noted that even the notion of *formal semantics* does not necessarily rely on model theory. Semantics based on order theory or on metric spaces, as used in denotational semantics of programming languages, are just one example, and can be ported to the knowledge representation realm [16].

[9] it might be more accurate to say: misconstrued

[10] See [3] for some alternatives to precision and recall in a Semantic Web context. We restrict our discussion to precision and recall simply because they are well established. We do not claim that there are no good or better alternatives: future research will have to determine this.

then we can also change our view on what form this semantics must take. Why would a shared set of inferences have to consist of conclusions that are held to be either completely true or completely false? Wouldn't it be reasonable to enforce a minimum (or maximum) degree of believe in certain statements? Or a degree of certainty? Or a degree of trust? This would amount to agent A and agent B establishing their semantic interoperability not by guaranteeing that B holds for eternally true all the consequences that follow from the statements communicated by A, but rather by guaranteering that B shares a degree of trust in all the sentences that are derivable from the sentences communicated by A.

A similar argument can be made for the handling of inconsistency. Shouldn't a semantics for "shared inference" be able to sort out inconsistencies and different perspectives on the fly? We know that classical model theory cannot deal with these issues. And what about default assumptions and the occurrence of exceptions to them? Classically, these lead to inconsistency, but in "shared inference" it should be dynamically resolvable.

While these perspectives, again, appear to be compatible with well-known knowledge representation approaches using, e.g., fuzzy or probabilistic logics [21, 31], paraconsistent reasoning [22], non-monotonic [7, 12,20,25], or mixed approaches [30], it is an open question whether they carry far enough for realistic use cases. While apparently promising as conceptual ideas, these logics have not yet been shown to be applicable in practice other than in simplified settings. How they could work on the open Semantic Web remains, to this date, unclear.

To us, it appears to be a reasonable perspective, that these issues need to be resolved, practically, in a different manner, as described below. Formal semantics, using non-classical logics, can probably still serve as a gold standard for evaluating inference system performances, but realistic data and applications will most likely force us to deviate from classical automated reasoning grounds for computing shared inferences.

5. Computing shared inferences

To summarize the train of thought we have laid out so far, we see that, in order to realize the interoperability required by the Semantic Web, we

- require shared ontologies which carry a formal semantics,
- formal semantics acts as a gold standard but does not need to be computed in a sound and complete way, and
- systems should be able to deal with noise, different perspectives, and uncertainty.

Traditionally, systems for computing inferences are based on logical proof theory and realize sound and complete algorithms on the assumption that input data is monolithic, noise-free, and conveys a single perspective on a situation or domain of applications. While this approach is certainly valid as such, it faces several severe challenges if ported to the Semantic Web. Two of the main obstacles are scalability of the algorithms, and requirements on the input data.

Concerning scalability, reasoning systems have made major leaps in the recent past [33,34]. However, it remains an open question when (and if[11]) these approaches will scale to the size of the web, and this problem is aggravated by the incorporation of nonclassical semantics as discussed in Section 4, which inherently brings a rapid decrease in efficiency.

Concerning requirements on the input data, it is quite unrealistic to expect that data from the open Semantic Web will ever be clean enough such that classical reasoning systems will be able to draw useful inferences from them. This would require Semantic Web data to be engineered strongly according to shared principles, which not only contrasts with the bottom-up nature of the Web, but is also unrealistic in terms of conceptual realizability: many statements are not true or false, they rather depend on the perspective taken.

If we come to the conclusion that inference systems based on logical proof theory likely will not work on web-scale realistic Semantic Web data,[12] the discussion from Section 3 becomes of central importance: Formal semantics is required as a gold standard for evaluation of systems computing shared inferences, however, it is okay for such systems to deviate from the gold standard, in a manner which can be qualitatively assessed in terms of precision and recall, if they scale better and/or are able to deal with realistic, noisy, data.

[11] Since the web keeps growing, they may never scale, even if they become much more efficient.

[12] This does, obviously, not preclude them from being very useful for smaller and/or more controled domains.

6. What is needed?

We have argued for the need of methods for computing shared inferences, which are not foremost based on the idea of producing sound and complete systems. We believe that there is a need for a concerted effort in the Semantic Web community to address this issue, both in terms of producing such systems, and in terms of pursuing use cases involving shared inference which employ reasoning methods which can scale up to web size.

Potential methods for establishing such inference systems can be found in other realms, where the need for approximate solutions is an accepted fact. Approximate algorithms, e.g., are commonly employed for NP-hard problems.[13] Approximate reasoning, understood in the same sense, has an established tradition. The development of according ideas for semantic web reasoning is indeed being pursued to a certain extent [18,26,27,32], and would benefit from a critical mass of further research.

Alternative approaches may employ methods which do not involve proof-theoretic aspects at all. From a bird's eye perspective, reasoning can be understood as a classification problem: classify a query as "true" or as "false." Machine learning, nature-inspired computing, or any method used in data mining or information retrieval are candidates for exploring new Semantic Web reasoning paradigms (see, e.g., [5,4,8,9,10]). These methods often have the pleasing property to be robust with respect to noise or contradictory input, and so there is reason to believe that they may simply render the difficulties identified in Section 4 to be void.

Let us close by emphasizing again that taking such approaches does not mean that we give up on formal semantics. It still serves as a gold standard for evaluation. It just means that we acknowledge that we need to rethink the role of semantics and the role of computation of semantics, provided we hope to make significant advances in the Semantic Web quest.

Acknowledgements. We thank Prateek Jain for digging out the example in Figure 1. Pascal Hitzler acknowledges support by the Write State University Research Council.

[13] Considering the fact that OWL reasoning is harder than NP, it is unfathomable why there should be any resistance against using approximate methods for OWL reasoning.

References

[1] S. Auer. Making the web a data washing machine – towards creating knowledge out of interlinked data. *Semantic Web – Interoperability, Usability, Applicability*, 2010. In this issue.

[2] C. Bizer, T. Heath, and T. Berners-Lee. Linked Data – the story so far. *International Journal on Semantic Web and Information Systems*, 5(3):1–22, 2009.

[3] C. d'Amato, N. Fanizzi, and F. Esposito. Query answering and ontology population: An inductive approach. In S. Bechhofer, M. Hauswirth, J. Hoffmann, and M. Koubarakis, editors, *The Semantic Web: Research and Applications, 5th European Semantic Web Conference, ESWC 2008, Tenerife, Canary Islands, Spain, June 1-5, 2008, Proceedings*, volume 5021 of *Lecture Notes in Computer Science*, pages 288–302. Springer, 2008.

[4] C. d'Amato, N. Fanizzi, and F. Esposito. Inductive learning for the Semantic Web: What does it buy? *Semantic Web – Interoperability, Usability, Applicability*, 2010. In this issue.

[5] C. d'Amato, N. Fanizzi, B. Fazzinga, G. Gottlob, and T. Lukasiewicz. Combining Semantic Web search with the power of inductive reasoning. In *Proceedings SUM 2010*, Lecture Notes in Computer Science. Springer, 2010. To appear.

[6] L. Ding, J. Shinavier, T. Finin, and D. L. McGuinness. An empirical study of owl:sameAs use in Linked Data. In *Proceeedings of the WebSci10: Extending the Frontiers of Society On-Line, April 2010, Raleigh, NC*, 2010. To appear.

[7] T. Eiter, G. Ianni, T. Lukasiewicz, R. Schindlauer, and H. Tompits. Combining Answer Set Programming with Description Logics for the Semantic Web. *Artificial Intelligence*, 172(12–13):1495–1539, August 2008.

[8] N. Fanizzi, C. d'Amato, and F. Esposito. Statistical learning for inductive query answering on OWL ontologies. In A. P. Sheth, S. Staab, M. Dean, M. Paolucci, D. Maynard, T. W. Finin, and K. Thirunarayan, editors, *The Semantic Web – ISWC 2008, 7th International Semantic Web Conference, ISWC 2008, Karlsruhe, Germany, October 26-30, 2008. Proceedings*, volume 5318 of *Lecture Notes in Computer Science*, pages 195–212. Springer, 2008.

[9] B. Fazzinga, G. Gianforme, G. Gottlob, and T. Lukasiewicz. Semantic Web search based on ontological conjunctive queries. In S. Link and H. Prade, editors, *Foundations of Information and Knowledge Systems, 6th International Symposium, FoIKS 2010, Sofia, Bulgaria, February 15-19, 2010. Proceedings*, volume 5956 of *Lecture Notes in Computer Science*, pages 153–172. Springer, 2010.

[10] B. Fazzinga and T. Lukasiewicz. Semantic Search on the Web. *Semantic Web – Interoperability, Usability, Applicability*, 2010. In this issue.

[11] D. Fensel and F. van Harmelen. Unifying reasoning and search to web scale. *IEEE Internet Computing*, 11(2):96, 94–95, 2007.

[12] S. Grimm and P. Hitzler. Semantic matchmaking of web resources with local closed-world reasoning. *International Journal of e-Commerce*, 12(2):89–126, 2008.

[13] H. Halpin and P. J. Hayes. When owl:sameAs isn't the same: An analysis of identity links on the Semantic Web. In *Proceedings of the WWW2010 workshop on Linked Data on the Web, LDOW2010*, 2010. To appear.

[14] P. Hayes, editor. *RDF Semantics*. W3C Recommendation, 10 February 2004. Available from

http://www.w3.org/TR/rdf-mt/.

[15] P. Hitzler, M. Krötzsch, B. Parsia, P. F. Patel-Schneider, and S. Rudolph, editors. *OWL 2 Web Ontology Language: Primer*. W3C Recommendation 27 October 2009, 2009. Available from http://www.w3.org/TR/owl2-primer/.

[16] P. Hitzler and A. K. Seda. *Mathematics Aspects of Logic Programming Semantics*. CRC Press, 2011. To appear.

[17] A. Hogan, A. Harth, A. Passant, S. Decker, and A. Polleres. Weaving the pedantic web. In *3rd International Workshop on Linked Data on the Web (LDOW2010) at WWW2010, Raleigh, USA, April 2010*, 2010. Available from http://events.linkeddata.org/ldow2010/.

[18] A. Hogan, A. Harth, and A. Polleres. Scalable authoritative OWL reasoning for the web. *Int. J. Semantic Web Inf. Syst.*, 5(2):49–90, 2009.

[19] P. Jain, P. Hitzler, P. Z. Yeh, K. Verma, and A. P. Sheth. Linked Data is Merely More Data. In D. Brickley, V. K. Chaudhri, H. Halpin, and D. McGuinness, editors, *Linked Data Meets Artificial Intelligence*, pages 82–86. AAAI Press, Menlo Park, CA, 2010.

[20] M. Knorr, J. J. Alferes, and P. Hitzler. A coherent well-founded model for Hybrid MKNF knowledge bases. In M. Ghallab, C. D. Spyropoulos, N. Fakotakis, and N. Avouris, editors, *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI-08)*, pages 99–103. IOS Press, 2008.

[21] T. Lukasiewicz and U. Straccia. Managing uncertainty and vagueness in description logics for the Semantic Web. *Journal of Web Semantics*, 6(4):291–308, 2008.

[22] Y. Ma and P. Hitzler. Paraconsistent reasoning for OWL 2. In A. Polleres and T. Swift, editors, *Web Reasoning and Rule Systems, Third International Conference, RR 2009, Chantilly, VA, USA, October 25-26, 2009, Proceedings*, volume 5837 of *Lecture Notes in Computer Science*, pages 197–211. Springer, 2009.

[23] F. Manola and E. Miller, editors. *Resource Description Framework (RDF). Primer*. W3C Recommendation, 10 February 2004. Available from http://www.w3.org/TR/rdf-primer/.

[24] B. Motik, P. F. Patel-Schneider, and B. C. Grau, editors. *OWL 2 Web Ontology Language: Direct Semantics*. W3C Recommendation, 27 October 2009. Available from http://www.w3.org/TR/owl2-direct-semantics/.

[25] B. Motik and R. Rosati. Reconciling description logics and rules. *Journal of the ACM*, 57(5), 2010.

[26] E. Oren, S. Kotoulas, G. Anadiotis, R. Siebes, A. ten Teije, and F. van Harmelen. Marvin: Distributed reasoning over large-scale Semantic Web data. *Web Semantics: Science, Services and Agents on the World Wide Web*, 7(4):305–316, 2009.

[27] J. Z. Pan and E. Thomas. Approximating OWL-DL ontologies. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, July 22-26, 2007, Vancouver, British Columbia, Canada*, pages 1434–1439. AAAI Press, 2007.

[28] A. Polleres, A. Hogan, A. Harth, and S. Decker. Can we ever catch up with the Web? *Semantic Web – Interoperability, Usability, Applicability*, 2010. In this issue.

[29] S. Rudolph, T. Tserendorj, and P. Hitzler. What is approximate reasoning? In D. Calvanese and G. Lausen, editors, *Web Reasoning and Rule Systems, Second International Conference, RR 2008, Karlsruhe, Germany, October 31-November 1, 2008. Proceedings*, volume 5341 of *Lecture Notes in Computer Science*, pages 150–164, 2008.

[30] T. Scharrenbach, R. Grütter, B. Waldvogel, and A. Bernstein. Structure preserving TBox repair using defaults. In V. Haarslev, D. Toman, and G. Weddell, editors, *Proc. of the 2010 Description Logic Workshop (DL 2010)*, volume 573 of *CEUR Workshop Proceedings*, 2010.

[31] G. Stoilos, G. B. Stamou, J. Z. Pan, V. Tzouvaras, and I. Horrocks. Reasoning with very expressive fuzzy description logics. *Journal of Artificial Intelligence Research*, 30:273–320, 2007.

[32] T. Tserendorj, S. Rudolph, M. Krötzsch, and P. Hitzler. Approximate OWL-reasoning with Screech. In D. Calvanese and G. Lausen, editors, *Web Reasoning and Rule Systems, Second International Conference, RR 2008, Karlsruhe, Germany, October 31-November 1, 2008. Proceedings*, volume 5341 of *Lecture Notes in Computer Science*, pages 165–180, 2008.

[33] J. Urbani, S. Kotoulas, J. Maassen, F. van Harmelen, and H. Bal. OWL reasoning with WebPIE: calculating the closure of 100 billion triples. In L. Aroyo, G. Antoniou, E. Hyvönen, A. ten Teije, H. Stuckenschmidt, L. Cabral, and T. Tudorache, editors, *The Semantic Web: Research and Applications, 7th Extended Semantic Web Conference, ESWC 2010, Heraklion, Crete, Greece, May 30 - June 3, 2010, Proceedings, Part I*, volume 6088 of *Lecture Notes in Computer Science*, pages 213–227. Springer, 2010.

[34] J. Urbani, S. Kotoulas, E. Oren, and F. van Harmelen. Scalable Distributed Reasoning Using MapReduce. In A. Bernstein et al., editors, *Proceedings of the 8th International Semantic Web Conference, ISWC 2009, Chantilly, VA, USA, October 25-29, 2009*, volume 5823 of *Lecture Notes in Computer Science*, pages 634–649. Springer, 2009.

[35] D. Vrandečić, M. Krötzsch, S. Rudolph, and U. Lösch. Leveraging non-lexical knowledge for the Linked Open Data web. In R. Héliot and A. Zimermann, editors, *5th Review of April Fool's day Transactions*, pages 18–27, 2010. Available from http://vmgal34.deri.ie/∼antzim/RAFT/afd2010.html.